arXiv:1406.4822v1 [cs.CG] 18 Jun 2014

Local Doubling Dimension of Point Sets

Aruni Choudhary* Michael Kerber†

June 19, 2014

Abstract

We introduce the notion of t-restricted doubling dimension of a point set in Euclidean space as the local intrinsic dimension up to scale t. In many applications information is only relevant for a fixed range of scales. We present an algorithm to construct a hierarchical net-tree up to scale t which we denote as the net-forest. We present a method based on *Locality Sensitive Hashing* to compute all near neighbours of points within a certain distance. Our construction of the net-forest is probabilistic, and we guarantee that with high probability, the net-forest is supplemented with the correct neighbouring information. We apply our net-forest construction scheme to create an approximate Čech complex up to a fixed scale; and its complexity depends on the local intrinsic dimension up to that scale.

1 Introduction

Motivation Often, one wants to perform tasks on data which lives in high dimensional spaces. Typically, algorithms for manipulating such high dimensional data take exponential time with respect to the ambient dimension. This is frequently quoted as the "curse of dimensionality". In many cases, however, practical input instances lie on low dimensional manifolds and a natural question arises as to how do we exploit this structural property for computationally feasible algorithms.

A well-established approach is to define a special notion of dimension on a point set. The *doubling dimension* of a point set P is the smallest integer Δ such that every ball centered at $p \in P$ of radius R is covered by at most 2^{Δ} non-empty balls of radius $R/2$ for any R. For instance, if P is a sample of an affine subspace of dimension k, it holds that $\Delta = k$, and often, $\Delta \ll d$ holds for more general samples of k-manifolds. A common goal is therefore to replace the exponential dependency on d by Δ in the complexity of geometric algorithms.

The concept of *(hierarchical) net-trees* can be seen as a generalization of quadtrees and allows for the translation of quadtree-based algorithms (which are exponential in d) to cases with small Δ. Technically, a net-tree provides a hierarchy of *nets* which summarize the point set in terms of a clustering scheme on different scales. For n points with doubling

*Max Planck Institute for Informatics, Saarbrücken, Germany (aruni.choudhary@mpi-inf.mpg.de)
†Max Planck Institute for Informatics, Saarbrücken, Germany (mkerber@mpi-inf.mpg.de)

dimension Δ, a net-tree can be constructed in expected $2^{O(\Delta)}O(n \log n)$ time, matching the time for constructing a quadtree except for replacing d with Δ [9]. As an application of particular importance, net-trees permit the efficient construction of *well-separated pair decomposition* (WSPD) which have various applications in geometric approximation, such as constructing *spanners*, finding *approximate nearest-neighbours*, approximating the *diameter* and the *closest-pair distance*.

In some applications, it is natural to upper bound the range of scales under investigation. In such cases, the doubling dimension does not capture the intrinsic complexity of the problem at hand, since it may be caused by a ball that is beyond the range of considered scales. Moreover, the net-tree construction of [9] proceeds in a top-down fashion, considering the high scales of the point set first. It therefore suffers from potentially bad large-scale properties of the point set, even when these properties are irrelevant for the given application.

Contributions In this paper, we introduce the concept of t-restricted doubling dimension Δ_t, which is the smallest integer such that any ball centered at $p \in P$ of radius $R \leq t$ is covered by at most 2^{Δ_t} non-empty balls of radius $R/2$. For simplicity of presentation, we restrict ourselves to the case of point sets in Euclidean space and postpone a more general treatment to an extended version of the paper. We present an algorithm to construct a *net-forest*, which contains the relevant data of a net-tree up to scale t. The runtime of the construction depends on Δ_{Ct} where C is a value independent of n and is defined later on. We hence remove the dependence on the doubling dimension Δ. The major geometric primitive of our algorithm is to find all neighbours of a point $p \in P$ with a distance of at most $\Theta(t)$. We propose an approach based on *Locality Sensitive Hashing* (LSH) from [5]. We have a trade-off between the size of C and the exponent of n in the complexity bound. The LSH based construction of the net-forest yields an expected runtime of $O\left(dn^{1+\rho} \log n(\log n + (14/\rho)^{\Delta_{7t/\rho}})\right)$ where $\rho \in (0, 1)$ is a parameter which can be chosen to be as small as desired. Comparing this bound with the full net-tree construction, our approach makes sense if $n^\rho \log n \ll 2^{O(\Delta)}$, that is, Δ is sufficiently large and $\Delta_{O(t)} \ll \Delta$.

As a consequence of our result, we can construct the part of the WSPD where all pairs are in distance at most $\Theta(t)$, adapting the construction scheme of [9, Sec.5]. That means that any application of WSPD that restricts its attention to low scales can profit from our approach.

As a further application, we show how to approximate Čech complexes using net-forests; Čech complexes are a standard tool for capturing topological properties of a point cloud. Such a complex depends on a scale parameter; in particular, in the context of persistent homology [6], Čech filtrations are considered, which encode Čech complexes at various scales. In [4], an approximate filtration of size $n(\frac{2}{\varepsilon})^{O(k \cdot \Delta)}$ has been constructed using net-trees. "Approximate" means that the two filtrations are interleaved in the sense of [3] and therefore yield similar persistence diagrams. However, because of the large size of filtrations, it is common to limit their construction to an upper threshold value t. With our results, we can construct such a upper-bound filtration of size $n(\frac{2}{\varepsilon})^{O(k \cdot \Delta_{O(t)})}$, thus replacing the doubling dimension in [4] by the $O(t)$-restricted doubling dimension.

Organization of the paper Section 2 gives a brief overview of doubling spaces and net-trees. We introduce the concept of the restricted version of the doubling dimension in Section 3. In Section 4 we present an algorithm to construct the net-forest up to a certain scale. Our algorithm uses the concept of LSH which we detail in Section 5. In Section 6 we give an overview of WSSDs and adapt their construction to use the net-forest. We summarize our results and conclude in Section 7.

2 Background

We fix P to be a finite point set consisting of n points throughout. As mentioned before, we restrict our attention to the Euclidean case $P \subset \mathbb{R}^d$, although some of the presented concepts could be extended to arbitrary metric spaces with some additional effort. In particular, the distance between any two points can be computed in $O(d)$ time for Euclidean setups.

Doubling dimension A *discrete ball* centered at a point $q \in P$ with radius r is the set of points $Q \subseteq P$ which satisfy $\|p - q\| \leq r$ for all $q \in Q$. The *doubling constant* [1, 12] is the smallest integer λ such for all $p \in P$ and all $r > 0$, the discrete ball centered at p of radius r is covered by λ discrete balls of radius $r/2$. The *doubling dimension* Δ of P is $\lceil \log_2 \lambda \rceil$. For example, a point set that is sampled from a k-dimensional subspace has a doubling dimension of k, independent of the ambient dimension d. In contrast, the d boundary points of the standard $(d-1)$-simplex form a doubling space of dimension $\lceil \log_2 d \rceil$. Even worse, we can construct a subset of doubling dimension $\Theta(d)$ by placing $2^{\Theta(d)}$ points inside the unit ball in $\mathbb{R}^d$ such that any two points have a distance of at least $3/2$ (the existence of such a point set follows by a simple volume argument). It is NP-hard to calculate the doubling dimension of a metric [8] but it can be approximated within a constant factor [9, Sec.9].

Nets and Net-trees A subset $Q \subseteq P$ is an (α, β)-*net*, denoted by $\mathcal{N}_{\alpha,\beta}$, if all points in P are in distance at most α from some point in Q and the distance between any two points in Q is at least β. Usually, α and β are coupled, that is, $\beta = \Theta(\alpha)$, in which case we talk about a *net at scale* α.

We can represent a nested sequence of nets for increasing scales α using a rooted tree structure, called the *net-tree* [9]. It has n leaves, each representing a point of P, and each internal node has at least two children. Every tree-node v represents the subsets of points given by the sub-tree rooted at v; we denote this set by P_v. Every v has a representative, $\text{rep}_v \in P_v$ that equals the representative of one of its children if v is not a leaf. Moreover, v is associated with an integer $\ell(v)$ called the *level* of v which satisfies $\ell(v) < \ell(\text{parent}(v))$, where $\text{parent}(v)$ is the parent of v in the tree. Finally, each node satisfies the following properties

- *Covering property*: $P_v \subseteq \mathbb{B}(\text{rep}_v, \frac{2\tau}{\tau-1} \cdot \tau^{\ell(v)})$
- *Packing property*: $P_v \supseteq P \bigcap \mathbb{B}(\text{rep}_v, \frac{\tau-5}{2\tau(\tau-1)} \cdot \tau^{\ell(\text{parent}(v))})$

where $\mathbb{B}(p, r)$ denotes the ball centered at p with radius r and $\tau = 11$.

The covering and packing properties ensure that each node v has at most $\lambda^{O(1)}$ children where λ is the doubling constant for P. Moreover, for any α, a net at scale α can be read off

from the net-tree immediately; see [9, Prop.2.2] for details. A net-tree can be constructed deterministically in time $2^{O(\Delta)}O(dn\log(n\cdot\Phi))$ where Φ represents the *spread* of P, using the greedy clustering scheme of Gonzalez [7] as a precursor to the tree construction. The dependence on spread can be eliminated by constructing the tree in $2^{O(\Delta)}dn\log n$ time in expectation (the additional factor of d compared to [9] accounts for the fact that we fixed the Euclidean metric, and therefore take into account the cost of computing distances in our computational model). The net-tree construction is oblivious to knowing the value of Δ. One can extract a net at scale ℓ [9, Pro.2.2] by collecting the set of nodes from T satisfying the condition $\mathcal{N}(\ell) = \{\mathrm{rep}_v | \ell(v) < \ell \leq \ell(\mathrm{parent}(v)\}$. The net-tree can be augmented to maintain, for each node u, a list of close-by nodes with similar diameter. Specifically, for each node u the data structure maintains the set

$$\mathrm{Rel}(u) := \{v \in N \mid \ell(v) \leq \ell(u) < \ell(\mathrm{parent}(v)) \textit{ and } \|\mathrm{rep}_u - \mathrm{rep}_v\| \leq 14\tau^{\ell(u)}\}. \qquad (1)$$

Rel(.) is computed during the construction without additional cost.

3 t-restricted doubling dimension

Definition 1. *The* t-restricted doubling constant *of* P *is the smallest positive integer* λ_t *such that all the points in any discrete ball centered at* $p \in P$ *of radius* r *with* $r \leq t$ *are covered by* λ_t *non empty balls of radius* $r/2$*. The corresponding* t-restricted doubling dimension Δ_t *is* $\lceil \log \lambda_t \rceil$.

By definition, $\Delta_t \leq \Delta$ for any P. More precisely, Δ_t is zero for t smaller than the closest-pair distance of P, and equals Δ when t is the diameter of P. While the doubling dimension of samples from an affine subspace of dimension k is bounded by k, this is not generally true for samples of k-manifolds where Δ increases due to curvature. To sketch an extreme example, consider an *almost space-filling* curve γ in $\mathbb{R}^d$ which has distance at most ε to any point of the unit ball, where ε is chosen small enough. We let P be a sufficiently dense sample of γ. While $\Delta_t = 1$ for small values of t, we claim that $\Delta_t = \Theta(d)$ for $\Delta_t = 1$; indeed, any sparser covering of the unit ball with balls of radius $1/2$ would leave some portion of the ball uncovered, and by construction, γ goes through that uncovered region, so that some point in P is missed. We skip a more formal treatment of this argument.

The "badness" of the previous example stems from the difference between Euclidean and geodesic distance of points lying on a lower-dimensional manifold. A common technique for approximating the geodesic distance is through the *shortest-path metric*: Let $G = (P, E)$ denote the graph whose edges are defined by the pairs of points of Euclidean distance at most t. The distance of two points p and q is then defined as the length of the shortest path from p to q (we assume for simplicity that G is connected). The concept of doubling dimensions extends to any metric space and we let Δ' denote the doubling dimension of P equipped with the shortest path metric. While Δ_t and Δ' appear to be related, Δ' can be much larger than Δ_t in general. Moreover, using the shortest-path metric raises the question of how to compute shortest path distances efficiently, if the cost of metric queries is taken into account.

4 Net-forests

We next define an appropriate data structure for point sets of small t-restricted doubling dimension, where t is a parameter of the construction. Informally, a *net-forest* is the subset of a net-tree obtained by truncating all nodes above scale t. More precisely, it is a collection of net-trees with roots $v_1, \ldots, v_k$ such that the representatives $\text{rep}_{v_1}, \ldots, \text{rep}_{v_k}$ form a (t,t)-net and the point sets $P_{v_1}, \ldots, P_{v_k}$ are disjoint and their union covers P. We define $\text{Rel}(u)$ for a node in the forest the same way as in (1) as the set of net-forest nodes that are close to u and of similar scale. As for net-trees, we call a net-forest *augmented* if each node u is equipped with $\text{Rel}(u)$.

Construction Our algorithm for constructing a net-forest is a simple adaption of the net-tree algorithm: we construct a (t,t)-net of P by clustering the point set and assign each point in P to its closest net-point. Each root in the net-forest represents one of the clusters. We also compute $\text{Rel}(u)$ for each root by finding the close-by clusters to u. Having this information, we can simply run the net-tree algorithm from [9] individually on each cluster to construct the net-forest. For augmenting it, we use the top-down traversal strategy as described in [9, Sec.3.4], inferring the neighbours of a node from the neighbours of its parent – since we have set up $\text{Rel}(\cdot)$ for the roots of the forest, this strategy is guaranteed to detect neighbours even if they belong to different trees of the forest.

Both the initial net construction and the $\text{Rel}(\cdot)$-construction require the following primitive for a point set Q, which we call a *near-neighbour query*: *Given a point $q \in Q$ and a radius r, return a list of points in Q containing exactly the points at distance r or smaller from q).* In the remainder of the section, we give more details next on how to compute net and the associated clusters, and how to find the neighbours for each such cluster, assuming that we have a primitive which can perform near-neighbour queries. In Section 5, we show the implementation of such an primitive.

Net construction We construct the net using a greedy scheme: For any input point, store a pointer $\text{N}(p)$ pointing to the net point assigned to point p. Initially, $\text{N}(p) \leftarrow NULL$ for all p. As long as there is a point p with $\text{N}(p) = NULL$, we set $\text{N}(p) \leftarrow p$ and query the near-neighbour primitive to get a list of points with distance at most t from p. For any point q in the list we update $\text{N}(q) \leftarrow p$ if either $\text{N}(q) = NULL$ or $\|p - q\| < \|N(q) - q\|$. Then we pick the next point p with $\text{N}(p) = NULL$.

At the end, the set of points p with $\text{N}(p) = p$ represent the net at scale t and points q satisfying $\text{N}(q) = p$ constitute $p's$ cluster. All points are assigned to their closest net-point. The net thus constructed is a (t,t)-net. Moreover, we assign the same level to all root clusters. In particular, for any root node v, we set $\ell(v)$ such that $\frac{2\tau}{\tau-1} \cdot \tau^{\ell(v)} = t$. Specifically, we set $\ell(v) = \lfloor \log_\tau \left(\frac{\tau-1}{2\tau} t\right) \rfloor$.

Computing the Rel(.) set for the roots After computing the net-points and their respective clusters, we need to augment the net-points with neighbouring information. Recall that $\text{Rel}(u)$ contains nodes in distance at most $14\tau^{\ell(u)}$ from rep_u. Since we have a (t,t)-net, the level of any root node u satisfies $14\tau^{\ell(u)} \leq 7t$. Hence we need to find neighbours of

net-points within $7t$, the minimum distance between any two net-points being more than t. By the doubling property, any root net-node can have at most $\lambda_{7t}^{\log_2 \frac{7t}{t/2}}$ such neighbours which simplifies to $C' = 14^{\Delta_{7t}}$. We use the near-neighbour primitive to compute such neighbours.

5 Near-neighbours primitive

We describe the primitive used in the previous section which performs near-neighbour queries. Our approach follows the notion of *Locality-sensitive hashing* (LSH) introduced by [10] for the Hamming metric and extended to Euclidean spaces in [5]. LSH is a popular approach to find approximate near-neighbours in high dimensions.

Locality Sensitive Hashing LSH applies several hash functions on a point set such that close points are more likely to map to the same hash-buckets than points which are sufficiently far away. A typical application of LSH is the (r, c)-*nearest neighbour* problem: If there exists a point within distance r of the query point q, report some point within distance cr of q, $c > 1$.

However, for our construction we wish to solve the following problem: report *all* points within distance r of the query point. We need the LSH oracle for two steps in our construction: constructing the net at scale t and computing the Rel$(\cdot)$ for the root-nodes. We show that both these steps requires a runtime sub-quadratic in n by a slight modification of the method presented in [5]. We repeat some of their definitions for clarity:

Definition 2. *A family of hash functions $\mathcal{H} = \{h : S \to U\}$ is called* (r_1, r_2, p_1, p_2)-sensitive *if for all $a, b \in S$, the following holds:*

- *if* $\|a - b\| \leq r_1, Pr_1 = P[h(a) = h(b)] \geq p_1$
- *if* $\|a - b\| \geq r_2, Pr_2 = P[h(a) = h(b)] \leq p_2$
- $p_1 \geq p_2$ *and* $r_1 \leq r_2$

We amplify the gap between Pr_1 and Pr_2 by concatenating k such hash functions, creating the family of hash functions $\mathcal{G} = \{g : S \to U^k\}$ such that $g(x) = (h_1(x), h_2(x), ..., h_k(x))$. For $g(x)$, we have the modified properties:

- if $\|a - b\| \leq r_1, P[g(a) = g(b)] \geq p_1^k$
- if $\|a - b\| \geq r_2, P[g(a) = g(b)] \leq p_2^k$

We describe our near-neighbour primitive next: The input is a point set Q with n points and a distance $r > 0$. As pre-processing step, we choose l hash functions $g_1, ..., g_l$ uniformly at random from $\mathcal{G}$ [5, Sec.3]. and hash each $p \in Q$ to the buckets $g_i(p) \forall i \in [1, l]$. Given a query point $q \in Q$, we iterate over $i = 1, \ldots, l$ and check for any point p in bucket $g_i(q)$ whether the distance to q is at most r. We output the points with this property as the near-neighbours of q.

We need to specify the parameters of LSH in the above description. Most importantly, we have to ensure that, with high probability, the output contains all points in distance r

from q. Moreover, we want the buckets to be of small size so that the primitive does not have to filter out too many false positives.

The performance of the LSH scheme depends upon a parameter $\rho \in (0,1)$ which appears as an exponent of n in the runtime. We choose the parameters p_1, p_2, r_1 and r_2 of the hashing scheme such that $\rho = \frac{\log p_1}{\log p_2} \approx \frac{r_1}{r_2}$ [5, Sec.4]. In the following parts of the section, we let $\rho = \frac{r_1}{r_2}$.

Lemma 1. *Let $r_1 := r$ and $r_2 := r/\rho$, $k := \lceil -\log_{p_2} n \rceil$ and $l := \lceil 2n^\rho \ln \frac{n}{\sqrt{\delta}} \rceil$ with an arbitrarily small constant δ. The near-neighbour primitive has the following properties:*

(i) With probability at least $1-\delta$, all points in distance at most r are reported for all query points.

(ii) For any query point q, the aggregate expected size of all buckets $g_1(q), ..., g_l(q)$ is at most $l(\widetilde{C}+1)$, where $\widetilde{C}$ is the number of points in Q with distance at most r_2 to q.

(iii) The pre-processing runtime is $O(dnkl)$ and the expected query runtime for a point is is $O(dl(k+\widetilde{C}))$, where $\widetilde{C}$ is defined as in (ii).

Proof. First we bound the expected aggregate size of the buckets. A bucket contains "close" points which are in distance at most r_2 from q and "far" points which are further away. However, since the probability of a far point falling in the same bucket as q is at most p_2^k, the expected size of a single bucket is at most $\widetilde{C} + np_2^k \leq \widetilde{C} + 1$ by our choice of k. Since there are l buckets, (ii) is satisfied.

For (i), fix two points $q_1, q_2 \in Q$ with distance at most r_1. We have to ensure that $g_j(q_1) = g_j(q_2)$ for some $j \in \{1, \ldots, l\}$; this implies that q_1 will be reported for query point q_2, and vice versa. The probability for $g_j(q_1) = g_j(q_2)$ for a fixed j is at least p_1^k, which is $p_1^{-\log_{p_2} n} = n^{-\rho}$. Hence the probability that $g_j(q_1) \neq g_j(q_2)$ holds for all $j \in \{1, \ldots, l\}$ is at most $(1-n^{-\rho})^l$ because we choose the hash functions uniformly at random. There are at most n^2 point pairs within distance at most r_1. By the union bound, the probability that at least one such pair maps into different buckets is at most $n^2(1-n^{-\rho})^l$. Now we can bound

$$\begin{aligned} n^2(1-n^{-\rho})^l &= n^2(1-n^{-\rho})^{2n^\rho \ln \frac{n}{\sqrt{\delta}}} \\ &= n^2(1-\frac{1}{n^\rho})^{n^\rho \ln \frac{n^2}{\delta}} \\ &\leq n^2 e^{-\ln \frac{n^2}{\delta}} = \delta, \end{aligned}$$

where we used the fact that $(1-1/x)^x \leq 1/e$ for all $x \geq 1$. It follows that the probability that all pairs of points in distance at most r_1 fall in at least one common bucket is at least $1-\delta$. This implies (i).

It remains to show (iii): in the pre-processing step, we have to compute $k \cdot l$ hash functions for n points. Computing the hash value for a point p, $h_i(p)$ takes $O(d)$ time [5, Sec.3.2]. For a query, we have to identify the buckets to consider in $O(dkl)$ time and then iterate through the (expected) $l(\widetilde{C}+1)$ candidates (using (ii)), spending $O(d)$ for each. □

Net-forest construction using LSH We analyze the complexity of our net-forest construction from Section 4 with the near-neighbour primitive that uses LSH. The primitive is used in the construction of the (t,t)-net, where we find the near-neighbours of distance at most t for a subset of points that form the net in the end. That means, we initialize the primitive with $r \leftarrow t$ and $Q \leftarrow P$.

Lemma 2. *The expected time to construct the (t,t)-net using LSH is*

$$O\bigg(dn^{1+\rho}\log n\Big(\log n + \Big(\frac{2}{\rho}\Big)^{\Delta_{t/\rho}}\Big)\bigg).$$

Proof. We consider the time spend on all near-neighbour queries: Let the resulting net consist of $m \leq n$ points. This implies that the algorithm proceeds in m rounds and queries the near-neighbours of m points. Let $\widetilde{C}_i$ denote the number of points in distance t/ρ from the i-th query point. By Lemma 1, the total complexity for all near-neighbour queries is:

$$O\left(ndkl + \sum_{i=1}^{m} dl(k+\widetilde{C}_i)\right) = O(ndkl + dl\sum_{i=1}^{m}\widetilde{C}_i) \tag{2}$$

We only need to bound the sum of the $\widetilde{C}_i$. For that, we fix some $q \in P$ and count in how many sets $\widetilde{C}_i$ may it appear. Let p_i denote the net-point chosen in the i-th iteration. We call such a net-point *close* to q if the distance to q is at most t/ρ. By definition, the net-points close to q lie in a ball of radius t/ρ centered at q. Since any pair of net-points has a distance of more than t, any ball of radius $t/2$ can contain at most one close net point. Following the definition of the t-restricted doubling dimension, the number of such net-points can be at most $\lambda_{t/\rho}^{\log_2 \frac{\frac{t}{\rho}}{\frac{t}{2}}}$ which simplifies to $\big(\frac{2}{\rho}\big)^{\Delta_{t/\rho}}$. It follows that

$$\sum_{i=1}^{m}\widetilde{C}_i \leq n\big(\frac{2}{\rho}\big)^{\Delta_{t/\rho}}.$$

Plugging this into (2), we get the claimed running time, observing that $k = O(\log n)$ and $l = O(n^{\rho}\log n)$ by Lemma 1. All additional operations in the net construction besides the calls of the primitive are dominated by that complexity. □

The second appearance of the near-neighbour primitive is in the construction of the Rel(.) sets for the roots of the net forest. Recall that the roots are represented by the net-points constructed before; let M denote the set of net-points and m their cardinality. We simply have to find all pairs of points of distance at most $7t$ among the net-points; and to do so we call the near neighbour primitive with $r \leftarrow 7t$ and $Q \leftarrow M$ for all $q \in M$.

Lemma 3. *Computing the* Rel(.) *sets using LSH takes expected time*

$$O\left(dn^{1+\rho}\log n\left(\log n + \left(\frac{14}{\rho}\right)^{\Delta_{7t/\rho}}\right)\right).$$

Proof. The proof is analogous to the proof of Lemma 2: Let $\widetilde{C}_i$ (for $i = 1, \ldots, m$) denote the number of net-points in distance at most $\frac{7t}{\rho}$ to the i-th net point. The same packing argument as in the previous Lemma shows that any $\widetilde{C}_i$ can be at most $(\frac{14}{\rho})^{\Delta_{7t/\rho}}$, so the their sum is bounded by $m(\frac{14}{\rho})^{\Delta_{7t/\rho}}$. Analogous to the proof of Lemma 2, we can thus bound the runtime to be as required, noting that $m \leq n$. □

Theorem 1. *The expected time for constructing the net-forest using LSH is*

$$O\left(dn^{1+\rho}\log n\left(\log n + \left(\frac{14}{\rho}\right)^{\Delta_{7t/\rho}}\right)\right).$$

Proof. Using Lemma 2 and Lemma 3, constructing the net and its Rel(·) sets are within the complexity bound. Constructing a single net-tree for a node containing n_i points takes time at most $2^{14\Delta_t} n_i \log n_i$ (the factor of 14 in $2's$ exponent can be seen by a careful analysis of [9, Sec.3.4]). Constructing individual net-trees for the clusters takes time: $\sum_{i=1}^{m} 2^{14\Delta_t} dn_i \log n_i$ Since $\sum_{i=1}^{m} n_i = n$, the above runtime simplifies to $2^{14\Delta_t} dn \log n$. Augmenting the net-forest takes time $dn2^{14\Delta_{7t}}$ [9, Sec.3.4]. The runtime for the latter steps are dominated by the Rel(·) and net construction for sufficiently large values of n. □

We see how the choice of ρ affects the complexity bound: For ρ very close to zero, we get a almost linear complexity in n, to the price that we have to consider larger balls in our algorithm and thus increase the restricted doubling dimension.

6 Applications

Well-Separated Pair Decomposition A pair of net-tree nodes (u, v) is *ε-well-separated* if $\max\{\text{diam}_u, \text{diam}_v\} \leq \varepsilon \text{dist}(u, v)$, where $\text{dist}(u, v)$ denotes the distance between the representatives of u and v. Informally speaking, all pairs of points (p, q) with $p \in P_u$, $q \in P_v$ have a similar distance to each other if (u, v) is well-separated. A *ε-well-separated pair decomposition* (ε-WSPD) is a collection of ε-well-separated pairs such that for any pair $(p, q) \in P \times P$, there exists a well-separated pair (u, v) such that $p \in P_u$ and $q \in P_v$; we say that such a pair (p, q) is *covered* by (u, v).

An ε-WSPD of size $n\varepsilon^{-O(\Delta)}$ can be computed in time $d\left(2^{O(\Delta)} n \log n + n(1/\varepsilon)^{O(\Delta)}\right)$ [9, Sec.5]. A WSPD considers pairs over all scales of distance, just because it has to cover any pair of points. We relax that condition and only require that all pairs of points in distance at most t are covered. We call the resulting structure *t-restricted ε-WSPD*.

We construct the t-restricted ε-WSPD as follows: We start by constructing the corresponding augmented net forest; let $u_1, \ldots, u_m$ be its roots. Since we know the Rel(·) set for any root, we can identify pairs (u_i, u_j) such that u_i is in $\text{Rel}(u_j)$ and vice versa (this also includes pairs where $u_i = u_j$). For any such pair, we call `genWSPD`(u_i, u_j) from [9, Sec.5], which simply traverses the sub-trees until it finds well-separated pairs. We output the union of all pairs generated in this way.

Theorem 2. *For $0 < \varepsilon < 1$ and $t > 0$, our algorithm computes a t-restricted ε-WSPD of size $n\varepsilon^{-O(\Delta_{7t})}$ in expected time*

$$NF + dn\varepsilon^{-O(\Delta_{7t})}.$$

where NF is the complexity for computing the net-forest from Theorem 1.

Proof. For correctness, any pair of nodes generated is ε-well-separated by definition. For the relaxed covering property, consider a pair (p, q) of points in distance at most t. There are roots u_1, u_2 in the net-forest with $p \in P_{u_1}$ and $q \in P_{u_2}$. Since the diameter of u_1 and u_2 is at most $2t$, the distance of rep_{u_1} and rep_{u_2} is at most $5t \leq 14\tau^{\ell(u_i)}$. Therefore, $u_2 \in \text{Rel}(u_1)$ (and vice versa), and there will be a pair generated that covers (p, q).

For the size bound, we can use the same charging argument as in [9, Sec.5]. We can additionally ensure by our construction that in all doubling arguments, the radius of the balls in question is at most $7t$ and therefore replace the doubling dimension by Δ_{7t} in the bound. The running time follows because the number of recursive calls of `genWSPD` is proportional to the output size, and we spend $O(d)$ time per recursion step. □

Well Separated Simplical Decomposition The concept of well-separated simplical decomposition (WSSDs) of point sets, introduced by Kerber and Sharathkumar [11] and extended to doubling spaces by Choudhary et al [4], generalizes the concept of WSPD to larger tuples. A $(k+1)$-tuple $(v_0, v_1, \ldots, v_k)$ is called *ε-well separated* if each v_i is a node of the net-tree and for any ball $\mathbb{B}$ which contains at least one point of each v_i, it holds that

$$v_0 \cup v_1 \cup .... \cup v_k \subseteq (1+\varepsilon)\mathbb{B}$$

where $(1+\varepsilon)\mathbb{B}$ is the ball with same center as $\mathbb{B}$ and radius multiplied by $(1+\varepsilon)$. An (ε, k)-WSSD is a set of ε-well-separated tuples of size up to $k+1$ such that any k-simplex is covered by some tuple. In [4], an (ε, k)-WSSD of size $n(2/\varepsilon)^{O(\Delta \cdot k)}$ is constructed in time $d\left(2^{O(\Delta)} n \log n + n(2/\varepsilon)^{O(\Delta \cdot k)}\right)$.

Similar as before, we define a *t-restricted (ε, k)-WSSD* to be a collection of ε-well-separated tuples such that each k-simplex that fits into a ball of radius t is covered by a tuple. The statement is equivalent to the condition that the radius of the smallest minimum enclosing ball containing points from each node of the tuple is at most t.

Theorem 3. *A* t-restricted (ε, k)-WSSD *of size $n(\frac{2}{\varepsilon})^{O(\Delta_{7t} \cdot k)}$ can be computed in time*

$$NF + nd(\frac{2}{\varepsilon})^{O(\Delta_{7t} \cdot k)},$$

where NF is the complexity for computing the net-forest from Theorem 1. Within the same time bound, we can construct a sequence of approximation complexes $(\mathcal{A}_\alpha)_{\alpha \in [0,t]}$ of size $n(\frac{2}{\varepsilon})^{O(\Delta_{7t} \cdot k)}$ whose persistence module is an $(1+\varepsilon)$-approximation (in the sense that the two modules are interleaved *[3]) of the* truncated Čech module *$(\mathcal{C}_\alpha)_{\alpha \in [0,t]}$.*

We defer the description of the algorithm to construct the *t-restricted (ε, k)-WSSD* and the proof of Theorem 3 to Appendix A.

Approximating the t-doubling dimension One can approximate Δ_t for any point set P up to a constant factor by constructing a net-forest T of scale t over P. Let x denote the maximum out-degree of any node in T. Then $\log x$ is a constant approximation of Δ_t. This follows from the arguments of [9, Sec.9].

7 Conclusion and future work

In this paper we presented an algorithm to construct a hierarchical net-forest up to a certain scale and applied it to the construction of WSPDs and approximate Čech complexes. One possible optimization we have ignored in our analysis is that the packing arguments we use are for the complete point set. However, during the Rel(·) construction, we work with the net-points which satisfy certain packing properties. Since these constitute a subset of the original point set, they may have an even lower doubling dimensions which we could exploit. Finding more applications for the t-restricted doubling dimension is another direction which we would like to look into.

Acknowledgements This research is supported by the Max Planck Center for Visual Computing and Communication.

References

[1] Patrice Assouad. Plongements Lipschitziens dans $\mathbb{R}^n$. *Bulletin de la Societe Mathematique de France*, 111:429–448, 1983.

[2] M. Bădoiu and K. Clarkson. Smaller core-sets for balls. In *Proc. 14th ACM-SIAM Symp. on Discr. Alg.*, pages 801–802, 2003.

[3] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, SCG '09, pages 237–246, 2009.

[4] Aruni Choudhary, Michael Kerber, and R. Sharathkumar. Approximate Cech complexes in low and high dimensions. `http://people.mpi-inf.mpg.de/~achoudha/Files/Papers/ApproximateCech.pdf`

[5] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, SCG '04, pages 253–262, 2004.

[6] Herbert Edelsbrunner and John Harer. *Computational Topology. An Introduction.* American Mathematical Society, 2010.

[7] Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theor. Comput. Sci.*, 38:293–306, 1985.

[8] Lee-Ad Gottlieb and Robert Krauthgamer. Proximity algorithms for nearly doubling spaces. *SIAM J. Discrete Math.*, 27(4):1759–1769, 2013.

[9] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low dimensional metrics, and their applications. In *SIAM J. Comput*, pages 150–158, 2005.

[10] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, 1998.

[11] Michael Kerber and R. Sharathkumar. Approximate Cech complexes in low and high dimensions. In *International Symposium on Algorithms and Computation*, pages 666–676, 2013.

[12] Kunal Talwar. Bypassing the embedding: Algorithms for low dimensional metrics. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, STOC '04, pages 281–290, 2004.

A Proof of Theorem 3

Construction of the t-restricted WSSD We describe the algorithm to construct the t-restricted (ε, k)-WSSD and prove its correctness and runtime. In this appendix, we will heavily rely on the notations, algorithms, and results presented in [4]. The algorithm proceeds iteratively; for $k = 1$, we construct a $(2t)$-restricted $\varepsilon/2$-WSPD using the algorithm from Section 6. To construct Γ_{k+1} from Γ_k, we iterate over the tuples $\gamma \in \Gamma_k$. We use the scheme of [4, Sec.3], computing an approximate meb of γ and then exploring ancestors of v_0 and their descendants at appropriate levels. The only complication arises when the algorithm requests for an ancestor higher than root of the tree of v_0. In such a case, our algorithm uses the root as the ancestor. In the following lemma, we will show that with this approach, we still cover all simplices with meb radius of at most t.

Lemma 4. *The algorithm computes a t-restricted (ε, k)-WSSD.*

Proof. We show by induction that with modified ancestor search, we still cover all simplices with meb radius at most t. For $k = 1$, the correctness of the algorithm follows from Theorem 2 in Section 6, Let Γ_{k-1} cover all $(k-1)$-simplices γ which satisfy $\mathrm{rad}(\gamma) \leq t$. Consider any k-simplex $\sigma = (m_0, \ldots, m_k)$ with $\mathrm{rad}(\sigma) \leq t$. From [4, Lem.9], there exists a point (say m_k) such that $m_k \in 2\mathrm{meb}(\sigma')$ where $\sigma' := \sigma \setminus \{m_k\}$ and $2\mathrm{meb}(\sigma')$ represents a ball with twice the radius and the same center as $\mathrm{meb}(\sigma')$. Since σ' is a $(k-1)$-simplex and $\mathrm{rad}(\sigma') \leq \mathrm{rad}(\sigma) \leq t$, it is covered by some k-tuple $\gamma = (v_0, \ldots, v_{k-1}) \in \Gamma_{k-1}$. To prove correctness, we show that when our algorithm reaches tuple γ, it produces a $(k+1)$-tuple (γ, x) such that $m_k \in P_x$ which implies that the simplex σ is covered by the $(k+1)$-tuple (γ, x).

When handling γ, the algorithm searches for an ancestor of v_0 at an appropriate scale. If this ancestor is found within the tree of v_0 in the net-forest, the arguments from [4, Lem.12] carry over to ensure that a suitable x is found. So let us assume that the algorithm chooses the root of the tree of the net-forest that v_0 lies in. Call that root node a_0. The algorithm considers all nodes in $\mathrm{Rel}(a_0)$ and creates new tuples with their descendants. Moreover, the net-forest contains a leaf representing the point m_k; let a' denote the root of its tree. It suffices to show that $a' \in \mathrm{Rel}(a_0)$. Since $\mathrm{rad}(\sigma) \leq t$, the distance of m_0 and m_k is at most $2t$. Moreover, the distance of m_0 to rep_{a_0} is at most t, because the representatives of the roots form a (t,t)-net. The same holds for m_k and a'. Using triangle inequality, the distance of rep_{a_0} and $\mathrm{rep}_{a'}$ is at most $4t$. This implies that $a' \in \mathrm{Rel}(a_0)$. □

Lemma 5. *The size of the computed t-restricted (ε, k)-WSSD Γ_k is $n(\frac{2}{\varepsilon})^{O(\Delta_{7t} \cdot k)}$.*

Proof. The proof of [4, Lem.13] carries over directly – indeed, we can replace all occurrences of Δ by Δ_{7t}. This comes from the fact that a node u has at most $14^{\Delta_{7t}}$ nodes in $\mathrm{Rel}(u)$, and for any node in $\mathrm{Rel}(u)$ we reach descendants of a level of at most $O(\log(2/\varepsilon))$ smaller then u (see the proof of [4, Lem.13] for details). Since every node in the net-forest has at most $2^{O(\Delta_t)}$ children, we create at most

$$14^{\Delta_{7t}}(\frac{2}{\varepsilon})^{O(\Delta_t)} = (\frac{2}{\varepsilon})^{O(\Delta_{7t})}$$

tuples in Γ_k from a tuple in Γ_{k-1}. With that, the bound can be proved by induction. □

Lemma 6. *Computing a t-restricted (ε)-WSSD takes expected time*

$$NF + nd(2/\varepsilon)^{O(\Delta_{7t} \cdot k)}$$

where NF is the complexity for computing the net-forest from Theorem 1.

Proof. The proof is analogous to [4, Lem.14], plugging in the running time for t-restricted ε-WSPD from Theorem 3 and the size bound from Lemma 5. □

Computing the approximate Čech filtration We use the scheme of [4, Sec.4] to construct the $(1+\varepsilon)$-approximate filtration on the t-restricted WSSD. The original construction works without modification. Using the notation from [4, Sec.4]., for any WST $\sigma = (v_0, v_1, \ldots, v_k)$ with $\ell(v_i) \leq h$, we add $\sigma' = (\mathrm{vcell}(v_0, h), \mathrm{vcell}(v_1, h), \ldots, \mathrm{vcell}(v_k, h))$ to $\mathcal{A}_\alpha$ if $\mathrm{rad}(\sigma') \leq \theta_\Delta$. The only potential problem with the t-restricted case is that such a vcell() might be a node higher than a root of the net-forest. This cannot happen, however, since h is chosen such that

$$\frac{2\tau}{\tau - 1}\tau^h \leq \frac{\varepsilon}{7}\alpha.$$

Since $\alpha \leq t$ and $\varepsilon \leq 1$, we have that

$$h < \lfloor \log_\tau \frac{\tau - 1}{2\tau} t \rfloor = \ell(u)$$

for any root u in the net-forest.